

5

TITLE OF THE INVENTION

BI-DIRECTIONAL WAVELENGTH SWITCHED RING OPTICAL
PROTECTION SWITCHING WAVELENGTH ASSIGNMENT ALGORITHM

CROSS REFERENCE TO RELATED APPLICATIONS

10

This application claims priority under 35 U.S.C.
§119(e) to provisional patent application serial number
60/216,892 filed July 7, 2000, the disclosure of which is
hereby incorporated by reference.

15

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR
DEVELOPMENT

N/A

20

BACKGROUND OF THE INVENTION

25

Wavelength Division Multiplexing (WDM) ring networks
are commonly known for use in commercial products and are
being widely used by telecommunication carriers. In
specific applications, WDM networks are used in the form
of rings and higher level networks are used in the form
of SONET/SDH self-healing rings. For these networks to
operate, it is necessary to provide switching equipment
in the nodes for terminating, originating and
regenerating the traffic in the lightpaths. This

-1-

switching equipment contributes significantly to the overall cost of the network.

One aspect of WDM ring networks is the assignment of wavelengths to carry the traffic over the lightpaths.

5 The problem of wavelength assignment for the lightpaths has been previously addressed by minimizing the number of wavelengths required for the given lightpath connections. However, by restricting the design analysis to minimizing the number of required wavelengths, the overall cost of the entire network may be neglected. A given wavelength assignment may require more switching equipment than alternative assignments using more wavelengths. Therefore, a method is desired for assigning wavelengths for a set of lightpath connections in which the overall cost of the system based on the number of assigned wavelengths and the switching equipment is considered.

BRIEF SUMMARY OF THE INVENTION

20 The present invention is directed to assigning the working path connections in a WDM ring network that minimizes the combined needs for wavelengths and Wavelength Interface Cards (WICs). The methods for assigning wavelengths to the working path connections according to the embodiments of the present invention reduces the number of required wavelengths to a value close to the optimal solution for solving the wavelength assignment problem alone and the number of required WICs to a value close to the optimal solution for solving the switching equipment problem alone. By solving these

problems in combination, the combined needs of wavelengths and WICs are minimized.

In one disclosed method, the wavelengths for the working path connections are assigned so that lightpaths are connected to form complete circles around the ring where possible. In another method, components of lightpaths are connected so that a linked component group is formed around the ring until no more connectable lightpath connections are available or until a complete circle is formed around the ring.

Other aspects, features and advantages of the present invention are disclosed in the detailed description that follows.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

The invention will be more fully understood by reference to the following detailed description of the invention in conjunction with the drawings, of which:

Figs. 1(a), 1(b), and 1(c) illustrate span switching, link switching, and path switching protection mechanisms, respectively, as known in the art;

Figs. 2(a), 2(b), and 2(c) illustrate bandwidth allocation utilizing path switching in a BWSR;

Figs. 3(a) and 3(b) illustrate the efficiency of wavelength usage in paired protection schemes for BWSRs;

Figs. 4(a) and 4(b) illustrate the efficiency of switching equipment usage in paired protection schemes for BWSRs;

Fig. 5 illustrates a BWSR utilizing a paired protection scheme;

Figs. 6(a) and 6(b) are flow charts illustrating light path grouping schemes; and

5 Figs. 7(a), 7(b), 7(c), and 7(d) illustrate working path assignments in paired protection schemes.

DETAILED DESCRIPTION OF THE INVENTION

10 Various protection architectures are known for Wavelength Division Multiplexing (WDM) in optical fiber ring networks. One such protection architecture is a Bidirectional Wavelength Switched Ring (BWSR) architecture that supports duplex traffic where working paths and protection paths are allocated for each
15 lightpath connection in the BWSR. Connections are typically established over the working path. However, when a network failure occurs in which the working path for a connection is no longer available, the connection is quickly switched to its protection path to restore
20 communication.

Examples of protection mechanisms that may be implemented on a BWSR include span switching 100, link switching 102, and path switching 104 as illustrated in Figs. 1(a), (b), and (c) for nodes 110, 120, 130, and
25 140. In each of Figs. 1(a), (b), and (c), communication is desired between nodes 110 and 120. In span switching 100, spare channels (such as wavelength or fiber, for example) are provisioned for every span 112, 122, 132, and 142 between the nodes. If a failure 170 occurs

between nodes 130 and 140, communication for a lightpath 150 in a clockwise direction between nodes 110 and 120 cannot be completed. In this case, span switching switches to a spare channel 160 for the span 132 which bypasses the failure 170 and provides communication between nodes 110 and 120. In link switching 102 as illustrated in Fig. 1(b), a lightpath protects each of the links 114, 124, 134, and 144 between the nodes. In this example a failure 172 occurs between nodes 130 and 140, and communication will be switched to lightpath 162.

Path switching 104 includes working paths and protection paths wherein a protection path is provided for every working path such that the protection path is edge-disjointed with the corresponding working path. For example, Fig. 1(c) illustrates a working path 154 and its respective protection path 164. When a failure 174 occurs in the working path 154 between the nodes 130 and 140 as shown in this example, the traffic is switched to the protection path 164 in a route completely different from the working path 154 such that the working path 154 is not traversed. In contrast with other switching techniques where switching equipment is required at every node that a connection traverses, the amount of necessary switching equipment is reduced in path switching because switching equipment is only required at the end nodes.

Figs. 2(a), (b), and (c) depict a BWSR 200 utilizing path switching and the bandwidth allocation for exemplary connections. The BWSR 200 includes a plurality of nodes 211, 212, 213, 214, 215, and 216 connected by optical

fiber links 221, 222, 223, 224, 225, and 226. Two wavelengths, λ_1 and λ_1' , are used for the working path and protection path connections, respectively. The working path connections are used for communication under normal conditions, and the protection path connections are used for communication when failures occur on the working paths. In Fig. 2(a), a first lightpath connection 230 originates at node 211 and terminates at node 213, and a second lightpath connection 232 originates at node 213 and terminates at node 215.

If a link fails anywhere within a working path connection, the connection is switched to the corresponding protection path. For instance, Fig. 2(b) illustrates the failure of the link 221 such that normal communication is broken over the wavelength λ_1 for the working path 230. As a result, switching mechanisms (not shown) at nodes 211 and 213 switch the connection to protection path connections 240 and 242 on wavelength λ_1' so that the connection still originates at node 211 and terminates at node 213. However, this protection path connection is in a different direction and on a different wavelength than the original connection. Similarly, if link 223 fails for the working path 232 as illustrated in Fig. 2(c), the switching mechanisms (not shown) at nodes 213 and 215 switch the connection to protection path connections 244 and 240 on wavelength λ_1' . Thereby, the protection path connection still originates at node 213 and terminates at node 215.

It is realized that the bandwidth in wavelength λ_1' may potentially be used by one of two different protection path connections (in Fig. 2(b), 240-242 and in Fig. 2(c), 244-240). In the case of a link failure, only one of the connections will use wavelength λ_1' . As long as only one link fails at a time, traffic from multiple working connections will never have to use the same protection channel. Therefore, more than one protection path may be assigned to the same wavelength λ_1' even while sharing some common links. Because of protection path bandwidth sharing, a protection path connection may undergo regenerations along the path that increases the amount of necessary network equipment. Specifically, as illustrated in Figs. 2(b) and 2(c), the protection path connection 240 is a segment that is shared in the protection paths for both working path connections 230 and 232. As a result, network equipment must regenerate the connection for these protection paths at node 215. Examples of the network equipment include Optical Filtering Cards (OFCs), and Wavelength Interface Cards (WICs) for performing add/drop, termination, regeneration, and optical pass-through for different wavelength channels at the nodes.

The network design problem for BWSR involves routing, grooming, wavelength assignment and network configuration. More particularly, given a set of traffic demands and network architecture, the equipment that should be installed at each node in the network and the network resources that should be allocated are

5 determined. The traffic demand includes a set of connections where each connection is defined by its source node, destination node, and requested bandwidth. The network architecture includes the number of nodes and the set of wavelengths supported by the network, the bandwidth capacity (in terms of OC-48 units, for example) for each wavelength, and the set of available OFCs.

10 A solution for the BWSR network design problem should consist of routing, wavelength, and timeslot assignment for each connection, and the configuration for each node. Specifically, the route (clockwise or counter-clockwise), the working and protection wavelengths assigned, and the timeslot(s) assigned (as a result of grooming) for each connection should be specified. Also, the set of OFCs to be installed at each node, and the WICs to be installed for each interface should be specified such that the routing and wavelength assignments for the connections are supported. For instance, the source and destination nodes should be equipped with an OFC that drops the wavelength and WICs terminating the wavelength should be installed at the appropriate node interfaces.

25 Paired protection is a protection wavelength assignment scheme having advantageous features for solving network design problems. In general, half of the designated wavelengths are dedicated for the working path connections and the other half are dedicated for the protection path connections. Furthermore, all of the working path connections that use a wavelength (λ , for

example) must use a complementary wavelength (λ' , in this example) for their protection path connections. One of the advantageous features of the paired protection scheme is that the network design problem is essentially reduced to wavelength assignment for the working path connections. Once the working path connections are assigned to a wavelength, the wavelength assignment for the protection path connections is automatically determined. Also, a working path wavelength is always protected by the same protection path wavelength so that paired protection is simple in terms of signaling and configuring. Furthermore, the structured relationship between the working and protection wavelengths simplifies the OFC architecture design because the OFC drops pairs of wavelengths.

However, the constraints associated with the structured relationship of the paired protection scheme may cause a sub-optimal wavelength assignment in terms of both the number of wavelengths and the number of WICs required in the network design. Figs. 3(a) and 3(b) illustrate an example where the paired protection scheme causes a sub-optimal usage of wavelengths. For the network 300 in Fig. 3(a), the connections between nodes 311 and 312 for a working path 320 and its protection path 330 are assigned to one wavelength, λ_1 . However, when the paired protection scheme is used for the same connections, one wavelength, λ_1 , is used for the working path connection and another wavelength, λ_2 , is used for the protection path connection as illustrated in Fig.

3(b). As a result, the connection for a protection path 332 is assigned to wavelength λ_2 while the working path 320 remains assigned to wavelength λ_1 and two wavelengths are sub-optimally used.

5 Figs. 4(a) and 4(b) illustrate an example where the paired protection scheme causes a sub-optimal usage of WICs. For the network 400 in Fig. 4(a), the connections between nodes 411 and 412 for a working path 420 and its protection path 430 are assigned to wavelength, λ_1 , and
10 the connections between nodes 413 and 414 for a working path 440 and its protection path 450 are assigned to wavelength, λ_2 . This results in the need of 8 WICS, 4 WICs for the working path connections (one WIC at each of the respective originating and terminating nodes 411,
15 412, 413, and 414) and 4 WICs for the protection path connections (another WIC at each of the respective originating and terminating nodes 411, 412, 413, and 414).

20 However, when the paired protection scheme is used for connections between the same nodes, one wavelength, λ_1 , is used for only the working path connections and the other wavelength, λ_2 , is used for only the protection path connections. As illustrated in Fig. 4(b), the working path connection 420 remains assigned to wavelength λ_1 but
25 a working path connection 442 between nodes 413 and 414 is now assigned to wavelength λ_2 . A protection path 432 for the working path 420 is then assigned to wavelength λ_2 while the protection path 450 remains assigned to wavelength λ_2 . In this case, a need for 12 WICS results,

4 WICs for the working path connections (one WIC at each of the respective originating and terminating nodes 411, 412, 413, and 414) and 8 WICs for the protection path connections (another WIC at each of the respective originating and terminating nodes 411, 412, 413, and 414 and an additional regeneration WIC at each of the nodes 411, 412, 413, and 414). Accordingly, the paired protection scheme sub-optimally uses 4 additional WICs.

In general, assuming that the wavelengths are assigned to connections, and k designates the number of connections assigned to each working wavelength, $4k$ WICs is the lower bound for the combined working and protection lightpaths for these k connections. For instance, when $k=1$, 4 WICs are needed for the connection and this is also the optimal configuration. When $k>1$, the connections need $2k$ WICs for the working path connections. With P representing the set of end nodes for the k connections, for each of the nodes in P , 2 WICs (1 WIC for the east connection and 1 WIC for the west connection) are needed for the protection path connections. Therefore, the number of WICs needed for the protection paths is $2P$. If $P=k$ (the connections cover a full circle), then the total number of WICs required is the optimal value of $4k$. The paired protection scheme may be more readily used if the combination of these optimal values for the needed WICs and wavelengths is minimized.

Fig. 5 illustrates an example of a solution to a network design problem for a BWSR 500 utilizing a paired

protection scheme. In this paired protection scheme, all connections on working paths having a wavelength λ_1 are protected by connections on protection paths having a wavelength λ_4 (the protection path connections are omitted for simplicity of the figure). Similarly, connections on working paths having wavelengths λ_2 and λ_3 are protected by connections on protection paths having wavelengths λ_5 and λ_6 , respectively. The BWSR includes nodes 511, 512, 513, 514, 515, and 516 supported by the wavelengths λ_1 , λ_2 , λ_3 , λ_4 , λ_5 , and λ_6 . The traffic demand is shown in Table 1 below:

Table 1

Connection	End Nodes	Bandwidth Requests (OC-48 units)
521	(1, 4)	4
522	(2, 3)	2
523	(2, 5)	2
524	(4, 6)	1
525	(5, 6)	1

Using shortest-path routing, the working wavelength assignment is shown in Table 2 below:

Table 2

Wavelength	Timeslot	Connection
λ_1	1	521
λ_1	2	521
λ_1	3	521
λ_1	4	521
λ_2	1	524
λ_2	2	525
λ_2	3	Idle
λ_2	4	Idle
λ_3	1	522
λ_3	2	522
λ_3	3	523
λ_3	4	523

5 Two types of optical filter cards (OFCs) are used in this example. A first OFC F_1 drops (λ_1 , λ_2 , λ_4 , and λ_5), and a second OFC F_2 drops (λ_2 , λ_3 , λ_5 , and λ_6). Thereby, the node configuration is shown in Table 3 below:

Table 3

Node	OFC
511	F_1
512	F_2
513	F_2
514	F_1
515	F_2
516	F_1

The primary objective in obtaining a solution is to find one routing, wavelength assignment and configuration that supports the traffic demand. A secondary objective is to minimize the cost of the network that is generally a function of the amount of required network equipment for the configuration. In the solution of Fig. 5, 20 WICs and 6 OFCs are required. It is generally desirable to find a solution for traffic demand that minimizes the amount of WICs and OFCs needed and reduces the cost of the network.

The methods for assigning wavelengths according to the embodiments of the present invention generally include grouping lightpaths and assigning wavelengths. In lightpath grouping, non-overlapping lightpaths are grouped so that each group of lightpaths is assigned to the same wavelength and the total number of required WICs is minimized.

In one embodiment of the present invention, a method of assigning the lightpath connections to complete a circle in paired protection schemes for network configurations is provided so that the required number wavelengths and WICs is minimized. If a set of connections exactly covers the BWSR, then the set of all of the end points are shared. By assigning the same wavelength to this set of connections, only 4k WICs are needed for the connections, which is the lower bound and optimal number of connections. The method according to the present embodiment examines a set of lightpaths and

finds a subset of the lightpaths that exactly covers the ring so that the usage of WICs is optimized.

Fig. 6(a) is a flow chart illustrating an embodiment of this method. At step 600, a set of lightpath
5 connections for the nodes in a ring is obtained. These connections represent the traffic demand. Next, the set of lightpath connections is examined at step 610 to determine whether a subset of lightpath connections exist which forms a circle. The subset of lightpath
10 connections will be determined to form a circle if at each node there is either one origination or termination of a lightpath connection or a single traversal of the node by a lightpath connection. If a determination is made at step 620 that a subset exists which forms a
15 circle, the subset is selected at step 630. Then, a unique wavelength is assigned to the subset at step 640. The selected subset is removed from the set at step 650 and the new set of lightpaths having the subset removed is examined again at step 610. Steps 610-650 are
20 repeated until no more subsets are determined to exist at step 620 and the method is completed at step 660.

In determining whether a subset of lightpath connections exists which form a circle at step 610, various algorithms may be used. In general, an algorithm
25 takes in a set of arcs and tries to divide the arcs into groups which forms a complete circle. An arc A is defined as a segment on a ring, which corresponds to a lightpath connection. A complete circle C is a set of arcs, which can be joined together to form a complete

circle. A partial circle Y is a set of arcs, which can be joined together to form a continuous part of a circle but not a complete circle. An arc x extends a partial circle Y if $Y \cup \{x\}$ is a partial circle or a complete circle.

In one embodiment of a find-circle algorithm, a RETRIEVE_CIRCLE module takes in a partial circle Y and a set of arcs A and recursively finds some arcs in A that completes the partial circle Y. In finding arcs in A to complete the partial circle Y, more than one arc may be returned as a potential candidate to extend the partial circle. Various criteria may be used for selecting from the arcs that extend the partial circle without forming the complete circle. For example, selecting the longest arc is one such selection criteria that may be used. Once a complete circle C is found, the set of arcs in the circle is returned. If no such set of arcs can be found, an empty set ϕ is returned.

In the main procedure of this algorithm, FIND_CIRCLES, a set of arcs A is input and the RETRIEVE_CIRCLE module is repeatedly called. Each time that the RETRIEVE_CIRCLE module returns a complete circle, FIND_CIRCLES saves the complete circle and removes those arcs from the input set of arcs. FIND_CIRCLES calls RETRIEVE_CIRCLE on the updated set of arcs repeatedly until RETRIEVE_CIRCLE returns an empty set ϕ . When the empty set is returned, no more circles can be found for the remaining arcs and the algorithm terminates.

7. if ($Y \neq \emptyset$) {
 8. return Y
 5 9. }
 10. }
 10 11. return \emptyset

In another embodiment of the present invention for minimizing the required wavelengths and WICs in paired
 15 protections schemes for network configurations, a method of assigning the lightpath connections for connected components is provided. In this method, a set of lightpaths sharing end nodes are made into a path as a connected component and all of the lightpaths in the set
 20 are assigned to the same wavelength.

Fig. 6(b) is a flow chart illustrating an embodiment of this method. At step 700, a set of lightpath connections for the nodes in the ring is obtained. One of the lightpath connections from the set is selected at
 25 step 710. This selected lightpath connection is removed from the set at step 720 to define a linked component group of one or more lightpath connections bounded by its originating and terminating nodes. At step 725, a determination is made as to whether any lightpath
 30 connections remain in the set. If at least one lightpath connection remains, the remaining lightpath connection(s) in the set are examined at step 730 for determining whether any of these connections either terminates at the

originating node of the linked component group or originates at the terminating node of the linked component group.

5 If it is determined at step 735 that one or more
lightpath connection(s) exist, the candidate(s) of
lightpath connection(s) are further analyzed at step 740
to determine whether they traverse any node(s) of already
connected lightpath(s). At step 745, if it is determined
10 that at least one of the candidate lightpath(s) does not
traverse any already connected lightpath(s), one
lightpath from the candidate lightpath(s) is selected at
step 750 and the process proceeds to step 760 for
examining any remaining lightpath connections. It is
15 appreciated that various rules may be applied for
selecting a lightpath when more than one lightpath
candidate exists. For example, one rule is to select the
longest linked component group that results in a minimum
number of connected components after its removal. Other
rules may be applied so that the lightpaths in the
20 connected components of the linked component group are
not fragmented after removing the set of selected
lightpaths.

25 If it is determined at any of steps 725, 735 or 745
that no more lightpath connections or candidates remain,
then the linked component group has been formed to its
longest extent possible. Accordingly, the defined linked
component group is assigned a unique wavelength at step
760. Similarly, if it is determined at step 765 that no
more lightpaths remain, the linked component group is

assigned a unique wavelength at step 770. Next, it is determined at step 770 whether any lightpath connections remain in the set. If at least one lightpath connection remains, the process proceeds to step 710 to select and examine the remaining lighpaths. Otherwise, the process ends at step 780 if no more lightpath connections remain in the set.

Figs. 7(a) and 7(b) illustrate examples of how the reduction of fragmented segments by the present method reduces the amount of needed switching equipment. In Fig. 7(a), two wavelengths λ_1 and λ_2 are provided for the four working path connections as shown in Table 4.

Table 4

Reference	Connection	Nodes	Wavelength
1	924	(1, 3)	λ_1
2	922	(4, 5)	λ_1
3	932	(2, 4)	λ_2
4	934	(5, 1)	λ_2

The protection path connections for Fig. 7(a) are shown in Table 5.

Table 5

Reference	Connection	Nodes	Wavelength
1'	948, 942, 944	(1, 5) (5, 4) (4, 3)	λ_1'
2'	944, 946, 948	(4, 3) (3, 1) (1, 5)	λ_1'
3'	956, 958, 952	(2, 1) (1, 5) (5, 4)	λ_2'
4'	952, 954, 956	(5, 4) (4, 2) (2, 1)	λ_2'

In this wavelength assignment example, 8 WICs (960-967) are needed for the working path connections and 16 WICS (970-985) are needed for the protection path connections for a total 16 WICs.

5 In the wavelength assignment of Fig. 7(b), connection 932 is shifted to the first wavelength λ_2 so that the working path connections are as shown in Table 6.

Table 6

Reference	Connection	Nodes	Wavelength
1	924	(1,3)	λ_1
2	922	(4,5)	λ_1
3	926	(5,1)	λ_1
4	932	(2,4)	λ_2

10

The corresponding protection path connections are as shown in Table 7.

Table 7

Reference	Connection	Nodes	Wavelength
1'	948, 942, 944	(1,5) (5,4) (4,3)	λ_1'
2'	944, 946, 948	(4,3) (3,1) (1,5)	λ_1'
3'	942, 944, 946	(5,4) (4,3) (3,1)	λ_1'
4'	950	(2,4)	λ_2'

15

In this wavelength assignment, 8 WICs (960-967) remain needed for the working path connections but only 10 WICS (970-977, 981, and 984) are needed for the

protection paths. Therefore, the wavelength assignment of Fig. 7(b) reduces the total number of WICs to 18, as compared to the 24 WICs needed in the wavelength assignment of Fig. 7(a). By moving the working path connection 934 on wavelength λ_2 in Fig. 7(a) to the working path connection 926 on wavelength λ_1 in Fig. 7(b), fragmented lightpath connections are eliminated so a linked component group of lightpaths exists. As a result, the amount of necessary switching equipment is reduced.

Follow-up processing may be carried out to remove the number of unshared connection end nodes in shared wavelengths. This follow-up processing is applicable for use after completion of all of the above-described methods including the methods for completing a circle and for assigning the lightpath connections for connected components. In one example, some of the isolated lightpaths are moved to unused wavelengths so that each unused wavelength will be assigned exactly one lightpath.

Figs. 7(c) and 7(d) illustrate how this follow-up processing can reduce the number of WICs in the network. In Fig. 7(c), all three lightpath connections 820, 822, and 824 for the working path connections are assigned to a first wavelength λ_1 while a second available wavelength λ_2 for the working path connections is not used. Segments of lightpath connections 830, 832, 834, 836 and 838 for the protection path connections on wavelength λ_1' are also shown. In this example, 6 WICs (one WIC at each of nodes 811, 813, 814 and 815 and two WICs at node 812) are

needed for the working path connections and 10 WICs (two WICs at each of nodes 811, 812, 813, 814 and 815) are needed for the protection path connections for a total of 16 WICs.

5 If the working path connection between nodes 814 and 815 is moved to the second wavelength λ_2 to form a working path connection 826, as illustrated in Fig. 7(d), the total number of WICs will be reduced by two. Segments of lightpath connections 840, 842, 844, and 850 for the
10 corresponding protection path connections on wavelengths λ_1' and λ_2' are also shown in Fig. 7(d). By shifting the one connection to the unused working path wavelength, only 8 WICs (one WIC at each of nodes 814 and 815 and two WICs at each of nodes 811, 812 and 813) are needed for
15 the protection path connections while 6 WICs remain necessary for the working path connections which leads to the reduction of two WICs.

 In assigning filters for each wavelength assignment input, a hitting set formulation is utilized. For each
20 node, a hitting set problem is formulated as a set of wavelengths that must be dropped (wavelengths assigned to circuits terminating at that node) and is defined as follows:

25 **INPUT:** A set of ground elements $G = \{e_1, e_2 \dots e_n\}$, each element e_i associated with weight w_i
 A family of sets $F = \{S_1, S_2 \dots S_m\}$, where each set S_i is a subset of G .

OUTPUT: A set $A \subseteq G$ that shares at least one
element with every S_i in F , i.e.,
 $\forall S_i \in F, A \cap S_i \neq \emptyset$.

OBJECTIVE: Minimize the weight of A , i.e., weight
sum of elements in A

In the hitting set formulation, weights are
associated with the filters to reflect the impact when
each filter is placed at a node. In one example of
associating a weight with a filter for a lightpath
assigned with a wavelength optically bypassing a node,
the weight of a filter dropping the wavelength at that
node is set to a large value. The weight is set to the
large value because an optical patch-through is required
for the wavelength at the node (and if an optical patch-
through is not allowed, the weight would be set to
infinity). The hitting set formulation provides an
assignment of filters in which the sum of the weights of
the filters is minimized. In general, a set of filters,
 F , with a minimum weight is sought so that the filter set
will drop all of the necessary wavelengths at each node.
The hitting-set formulation initializes a hitting-set
problem instance for each node, finds the wavelength
assigned, and updates the hitting-set instances for each
node for each lightpath p .

More particularly, the weight of the filters is
initialized and the set F is set to 0 before a lightpath
is set up for any node. When a lightpath p is assigned a
wavelength, λ_1 for example, the weight of each filter in

the node is updated. For instance, if the lightpath p bypasses a node i, then the weights for the filters that drop λ_1 are increased in the hitting-set formulation for node i. Additionally, the filter set F for the nodes where the lightpath is added or dropped is updated. The end points of the lightpath and intermediate nodes where circuits in the lightpath terminate are also updated. The filter set F for each of these nodes will include a new element S, which contains the set of filters that drop λ_1 .

Accordingly, the present methods for assigning working path connections in paired protection schemes minimizes the amount of switching equipment needed in network configurations. In particular, the wavelength connection assignment in the present invention reduces the combination of WICs and assigned wavelengths for the desired lightpath connections so that the advantages of the paired protection scheme may be realized while reducing the costs of the network.

It will be apparent to those skilled in the art that other modifications to and variations of the above-described techniques are possible without departing from the inventive concepts disclosed herein. Accordingly, the invention should be viewed as limited solely by the scope and spirit of the appended claims.